

CLAIMS

1. A method of allocating bandwidth at a network element, the method comprising the steps of:

metering first traffic for a first Per Hop Basis group (PHB) to ascertain first in-profile traffic for the first PHB;

metering second traffic for a second PHB to ascertain second in-profile traffic for the second PHB; and

commonly metering first traffic that has not been ascertained to be first in-profile traffic with second traffic that has not been ascertained to be second in-profile traffic to ascertain commonly metered traffic.

2. The method of claim 1, wherein the bandwidth allocated by the method is bandwidth on one or more links connected to a port of the network element, and wherein the first traffic and second traffic are configured to be transmitted through the port.

3. The method of claim 2, wherein the bandwidth is allocated to the first in-profile traffic, the second in-profile traffic, and wherein any surplus bandwidth not consumed by the first in-profile traffic and second in-profile traffic is allocated to the commonly metered traffic.

4. The method of claim 1, further comprising the step of classifying incoming traffic into the first traffic and the second traffic.

5. The method of claim 1, further comprising marking the first in-profile traffic with a first designation, marking the second in-profile traffic with the first designation, and marking at least a first portion of the commonly metered traffic with a second designation.

6. The method of claim 5, further comprising marking at least a second portion of the commonly metered traffic with a third designation.

7. The method of claim 1, wherein the step of metering the first traffic is performed using a first token bucket and wherein the ascertained first in-profile traffic for the first PHB is a portion of the first traffic for which there is sufficient tokens in the first token bucket;

wherein the step of metering the second traffic is performed using a second token bucket and wherein the ascertained second in-profile traffic for the second PHB is a portion of the second traffic for which there is sufficient tokens in the second token bucket; and

wherein the step of commonly metering is performed using a third common token bucket and wherein the ascertained commonly metered traffic is a portion of the first traffic that has not been ascertained to be first in-profile traffic and second traffic that has not been ascertained to be second in-profile traffic for which there is sufficient tokens in the third common token bucket.

8. The method of claim 7, wherein the first token bucket is provided with no tokens so that the ascertained first in-profile traffic for the first PHB is set to zero.

9. The method of claim 7, further comprising the steps of:

coloring green the portion of the first traffic for which there is sufficient tokens in the first token bucket;

coloring green the portion of the second traffic for which there is sufficient tokens in the second token bucket; and

coloring yellow the portion of the commonly metered traffic for which there is sufficient tokens in the third common token bucket.

10. The method of claim 9, further comprising the step of:

coloring red a portion of the commonly metered traffic for which there is not sufficient tokens in the third common token bucket.

11. A packet meter, comprising:

a Per Hop Behavior (PHB) classifier; and

a meter configured to meter in-profile packets on a PHB basis and out-of-profile packets on a common basis.

12. The packet meter of claim 11, further comprising a marker configured to mark packets metered by the meter.
13. The packet meter of claim 11, wherein the meter is configured to apply bandwidth allocation rules on a PHB basis to allocate bandwidth to each PHB based on its associated bandwidth allocation rule.
14. The packet meter of claim 13, wherein the bandwidth allocation rules comprise at least a committed information rate and a committed burst rate.
15. The packet meter of claim 13, wherein the meter is configured to meter out-of-profile packets on a common basis by allocating surplus bandwidth.
16. The packet meter of claim 11, wherein the meter comprises:
 - a first plurality of token buckets, said first plurality of token buckets being configured to allocate bandwidth to PHBs on a per-PHB basis; and
 - a common token bucket configured to meter out-of-profile packets on a per-port basis.
17. The packet meter of claim 16, further comprising a marker, said marker being configured to mark as green any packet that was passed by one of the first plurality of token buckets, and said marker being configured to mark as yellow any packet that was passed by the common token bucket.
18. The packet meter of claim 17, wherein the marker is further configured to mark as red any packet that was not passed by one of the first plurality of token buckets and not passed by the common token bucket.